

Reproducing Kernel Hilbert Spaces for Penalized Regression: A Tutorial.

Alvaro Nosedal-Sanchez^a, Curtis B. Storlie^b, Thomas C.M. Lee^c, Ronald Christensen^d

^a Indiana University of Pennsylvania

^b Los Alamos National Laboratory

^c University of California, Davis

^d University of New Mexico

Date: March 18, 2012

Abstract

Penalized Regression procedures have become very popular ways to estimate complicated functions. The smoothing spline, for example, is the solution of a minimization problem in a functional space. If such a minimization problem is posed on a reproducing kernel Hilbert space (RKHS), the solution is guaranteed to exist, is unique and has a very simple form. There are excellent books and articles about RKHS and their applications in statistics, however, this existing literature is very dense. This paper provides a friendly reference for a reader approaching this subject for the first time. It begins with a simple problem, a system of linear equations, then gives an intuitive motivation for reproducing kernels. Armed with the intuition gained from our first examples, we take the reader from Vector Spaces, to Banach Spaces and to RKHS. Finally, we present some statistical estimation problems that can be solved using the mathematical machinery discussed. After reading this tutorial, the reader will be ready to study more advanced texts and articles about the subject such as Wahba (1990) or Gu (2002).

1 Introduction

Penalized regression procedures have become a very popular approach to estimating complex functions (Wahba 1990, Eubank 1999, Hastie, Tibshirani & Friedman 2001). They are commonly used in areas such as functional data analysis (Ramsay & Silverman 2005), computer model analysis (Storlie, Swiler, Helton & Sallaberry 2009), image processing (Berman 1994), and various applications of spatial statistics (Bivand, Pebesma & Gómez-Rubio 2008), to name a few. Penalized regression procedures use an estimator that is defined as the solution

to a minimization problem. In any minimization problem, there are the following questions: Does the solution exist? If yes, is the solution unique? How can we find it? If the problem is posed in the Reproducing Kernel Hilbert Space (RKHS) framework that we discuss below, then the solution is guaranteed to exist, it is unique, and it takes a particularly simple form.

Reproducing Kernel Hilbert Spaces and Reproducing Kernels play a central role in Penalized Regression. The purpose of this article is to provide a constructive tutorial for statisticians interested in learning about RKHS methods in regression before studying more advanced texts and articles about the subject such as Wahba (1990), Gu (2002), or Pearce & Wand (2006). Pearce & Wand (2006) provide a review of the connection between penalized splines and support vector machines (SVMs) using the RKHS framework. This paper is intended to complement that review by providing the reader with the necessary background about RKHSs to fully understand how RKHS results are used in penalized regression problems. Much of the associated literature begins with picking a reproducing kernel and goes from there. Reproducing kernels may be the beginning of an application but they are the end of a body of theory. This paper explicates that body of theory in an effort to make its application to penalized regression (and hence SVMs) more lucid. We describe how to construct a kernel with the properties needed for a given application and how to use the properties of that kernel for penalized regression. We provide several examples to help motivate and solidify the concepts as well as a transparent justification for the so called “kernel trick”.

In the first section of this tutorial we present, and solve, simple problems while gently introducing key concepts. The remainder of the paper is organized as follows. Section 2 takes the reader from a basic understanding of fields through Banach Spaces and Hilbert Spaces. In Section 3, we provide elementary theory of RKHSs along with some examples. Section 4 discusses Penalized Regression with RKHSs. Two specific examples involving ridge regression and smoothing splines are given with code written in the R language (R Development Core

Team 2005) to solidify the concepts. Some methods for smoothing parameter selection are briefly mentioned. Section 5 contains some closing remarks.

1.1 Why we care about Reproducing Kernel Hilbert Spaces.

Before introducing new concepts, we present some simple illustrations of the tools used to solve problems in the RKHS framework. Consider solving the system of linear equations

$$x_1 + x_3 = 0 \tag{1}$$

$$x_2 = 1. \tag{2}$$

Clearly, the real-valued solutions to this system are the vectors $\mathbf{x}_*^t = (-\alpha, 1, \alpha)$ for $\alpha \in \mathfrak{R}$. Suppose we want to find the “smallest” solution. Under the usual squared norm $\|\mathbf{x}\|^2 = x_1^2 + x_2^2 + x_3^2$, the smallest solution is $\mathbf{x}_s^t = (0, 1, 0)$.

Now consider a more general problem. For a given $p \times n$ matrix \mathbf{R} and $n \times 1$ matrix $\boldsymbol{\eta}$, solve

$$\mathbf{R}^t \mathbf{x} = \boldsymbol{\eta}, \tag{3}$$

where \mathbf{R}^t is the transpose of \mathbf{R} , \mathbf{x} and the columns of \mathbf{R} , say \mathbf{R}_k , $k = 1, 2, \dots, n$, are all in \mathfrak{R}^p , and $\boldsymbol{\eta} \in \mathfrak{R}^n$. We wish to find the solution \mathbf{x}_s that minimizes the norm $\|\mathbf{x}\| = \sqrt{\mathbf{x}^t \mathbf{x}}$. We solve the problem using concepts that extend to RKHS.

A solution \mathbf{x}_* (not necessarily a minimum norm solution) exists whenever $\boldsymbol{\eta} \in C(\mathbf{R}^t)$. Here $C(\mathbf{R}^t)$ denotes the column space of \mathbf{R}^t . Given one solution \mathbf{x}_* , all solutions \mathbf{x} must satisfy

$$\mathbf{R}^t \mathbf{x} = \mathbf{R}^t \mathbf{x}_*$$

or

$$\mathbf{R}^t (\mathbf{x}_* - \mathbf{x}) = \mathbf{0}.$$

The vector \mathbf{x}_* can be written uniquely as $\mathbf{x}_* = \mathbf{x}_0 + \mathbf{x}_1$ with $\mathbf{x}_0 \in C(\mathbf{R})$ and $\mathbf{x}_1 \in C(\mathbf{R})^\perp$, where $C(\mathbf{R})^\perp$ is the orthogonal complement of $C(\mathbf{R})$ (i.e., $\mathbf{x}_0^t \mathbf{x}_1 = 0$, with orthogonality defined more formally on page 12). Clearly, \mathbf{x}_0 is a solution because $\mathbf{R}^t (\mathbf{x}_* - \mathbf{x}_0) = \mathbf{R}^t \mathbf{x}_1 = \mathbf{0}$

In fact, \mathbf{x}_0 is both the unique solution in $C(\mathbf{R})$ and the minimum norm solution. If \mathbf{x} is any other solution in $C(\mathbf{R})$ then $\mathbf{R}^t(\mathbf{x} - \mathbf{x}_0) = \mathbf{0}$ so we have both $(\mathbf{x} - \mathbf{x}_0) \in C(\mathbf{R})^\perp$ and $(\mathbf{x} - \mathbf{x}_0) \in C(\mathbf{R})$, two sets whose intersection is only the $\mathbf{0}$ vector. Thus $\mathbf{x} - \mathbf{x}_0 = \mathbf{0}$ and $\mathbf{x} = \mathbf{x}_0$. In other words, every solution \mathbf{x}_* has the same \mathbf{x}_0 vector. Finally, \mathbf{x}_0 is also the minimum norm solution because the arbitrary solution \mathbf{x}_* has

$$\mathbf{x}_0^t \mathbf{x}_0 \leq \mathbf{x}_0^t \mathbf{x}_0 + \mathbf{x}_1^t \mathbf{x}_1 = \mathbf{x}_*^t \mathbf{x}_*.$$

We have established the existence of a unique, minimum norm solution in $C(\mathbf{R})$ that can be written as

$$\mathbf{x}_s \equiv \mathbf{x}_0 = \mathbf{R}\boldsymbol{\xi} = \sum_{k=1}^n \xi_k \mathbf{R}_k, \quad (4)$$

for some ξ_k , $k = 1, \dots, n$. To find \mathbf{x}_s explicitly, write $\mathbf{x}_s = \mathbf{R}\boldsymbol{\xi}$ and the defining equation (3) becomes

$$\mathbf{R}^t \mathbf{R} \boldsymbol{\xi} = \boldsymbol{\eta}, \quad (5)$$

which is just a system of linear equations. Even if there exist multiple solutions $\boldsymbol{\xi}$, $\mathbf{R}\boldsymbol{\xi}$ is unique.

Now we use this framework to find the “smallest” solution to the system of equations (1) and (2). In the general framework we have

$$\mathbf{x}^t = (x_1, x_2, x_3),$$

$$\boldsymbol{\eta}^t = (0, 1),$$

$$\mathbf{R}_1^t = (1, 0, 1),$$

$$\mathbf{R}_2^t = (0, 1, 0).$$

We know that the solution has the form (4) and we also know that we have to solve a system of equations given by (5). In this case, the system of equations is

$$2\xi_1 + 0\xi_2 = 0,$$

$$0\xi_1 + 1\xi_2 = 1.$$

The solution to the system is $(\xi_1, \xi_2) = (0, 1)$ which implies that our solution to the original problem is $\mathbf{x}_s = 0\mathbf{R}_1 + 1\mathbf{R}_2 = (0, 1, 0)^t$ as expected.

Virtually the same methods can be used to solve a similar problem in any inner-product space Ω . As discussed later, an inner product $\langle \cdot, \cdot \rangle$ assigns real numbers to pairs of “vectors.” For given vectors $\mathbf{R}_k \in \Omega$ and numbers $\eta_k \in \Re$, find $\mathbf{x} \in \Omega$ such that

$$\langle \mathbf{R}_k, \mathbf{x} \rangle = \eta_k, \quad k = 1, 2, \dots, n \quad (6)$$

for which the norm of $\|\mathbf{x}\| \equiv \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ is minimal. The solution has the form

$$\mathbf{x}_s = \sum_{k=1}^n \xi_k \mathbf{R}_k, \quad (7)$$

with ξ_k satisfying the linear equations

$$\sum_{k=1}^n \langle \mathbf{R}_i, \mathbf{R}_k \rangle \xi_k = \eta_i, \quad i = 1, \dots, n.$$

For a formal proof see Máté (1990), pg. 70. In RKHS applications, vectors are typically functions. We now apply this result to the interpolating spline problem.

1.2 Interpolating Splines.

Suppose we want to find a function $f(t)$ that interpolates between the points (t_k, η_k) , $k = 0, 1, 2, \dots, n$, where $\eta_0 \equiv 0$ and $0 = t_0 < t_1 < \dots < t_n = 1$. We restrict attention to functions $f \in F$ where $F = \{f: f \text{ is absolutely continuous on } [0,1], f(0) = 0, \int_0^1 [f'(t)]^2 dt < \infty\}$. Throughout $f^{(m)}$ denotes the m -th derivative of f with $f' \equiv f^{(1)}$ and $f'' \equiv f^{(2)}$. The restriction that $\eta_0 = f(0) = 0$ is not really necessary, but simplifies the presentation.

We want to find the smoothest function $f(t)$ that satisfies $f(t_k) = \eta_k$, $k = 1, \dots, n$. Defining an inner product on F by

$$\langle f, g \rangle = \int_0^1 f'(x)g'(x)dx$$

implies a norm over the space F that is small for “smooth” functions. To address the interpolation problem, note that the functions $R_k(s) \equiv \min(s, t_k)$, $k = 1, 2, \dots, n$ have $R_k(0) = 0$ and the property that $\langle R_k, f \rangle = f(t_k)$ because

$$\begin{aligned}\langle f, R_k \rangle &= \int_0^1 f'(s) R'_k(s) ds \\ &= \int_0^{t_k} f'(s) 1 ds + \int_{t_k}^1 f'(s) 0 ds \\ &= \int_0^{t_k} f'(s) ds = f(t_k) - f(0) = f(t_k).\end{aligned}$$

Thus, an interpolator f satisfies a system of equations like (6), namely

$$f(t_k) = \langle R_k, f \rangle = \eta_k, \quad k = 1, \dots, n. \quad (8)$$

and by (7), the smoothest function f (minimum norm) that satisfies the requirements has the form

$$\hat{f}(t) = \sum_{k=1}^n \xi_k R_k(t)$$

The ξ_j 's are the solutions to the system of real linear equations obtained by substituting \hat{f} into (8),

$$\sum_{j=1}^n \langle R_k, R_j \rangle \xi_j = \eta_k, \quad k = 1, 2, \dots, n.$$

Note that

$$\langle R_k, R_j \rangle = R_j(t_k) = R_k(t_j) = \min(t_k, t_j)$$

and define the function

$$R(s, t) = \min(s, t)$$

which turns out to be a reproducing kernel.

Numerical example. Given points $f(t_i) = \eta_i$, say, $f(0) = 0$, $f(0.1) = 0.1$, $f(0.25) = 1$, $f(0.5) = 2$, $f(0.75) = 1.5$, and $f(1) = 1.75$, we find

$$\arg \min_{f \in F} \|f\|^2 = \int_0^1 f'(x)^2 dx.$$

The system of equations is

$$\begin{aligned} 0.1\xi_1 + 0.1\xi_2 + 0.1\xi_3 + 0.1\xi_4 + 0.1\xi_5 &= 0.1 \\ 0.1\xi_1 + 0.25\xi_2 + 0.25\xi_3 + 0.25\xi_4 + 0.25\xi_5 &= 1 \\ 0.1\xi_1 + 0.25\xi_2 + 0.5\xi_3 + 0.5\xi_4 + 0.5\xi_5 &= 2 \\ 0.1\xi_1 + 0.25\xi_2 + 0.5\xi_3 + 0.75\xi_4 + 0.75\xi_5 &= 1.5 \\ 0.1\xi_1 + 0.25\xi_2 + 0.5\xi_3 + 0.75\xi_4 + \xi_5 &= 1.75. \end{aligned}$$

The solution is $\xi = (-5, 2, 6, -3, 1)^t$, which implies that our function is

$$\hat{f}(t) = -5R_1(t) + 2R_2(t) + 6R_3(t) - 3R_4(t) + 1R_5(t) \quad (9)$$

$$= -5R(t, t_1) + 2R(t, t_2) + 6R(t, t_3) - 3R(t, t_4) + 1R(t, t_5) \quad (10)$$

or, adding the slopes for $t > t_i$ and finding the intercepts,

$$\hat{f}(t) = \begin{cases} t & 0 \leq t \leq 0.1 \\ 6t - 0.5 & 0.1 \leq t \leq 0.25 \\ 4t & 0.25 \leq t \leq 0.5 \\ -2t + 3 & 0.5 \leq t \leq 0.75 \\ t + 0.75 & 0.75 \leq t \leq 1 \end{cases}$$

This is the linear interpolating spline as can be seen graphically in Figure 1.

For this illustration we restricted f so that $f(0) = 0$. This was only for convenience of presentation. It can be shown that the form of the solution remains the same with any shift to the function, so that in general the solution takes the form $\hat{f}(t) = \xi_0 + \sum_{j=1}^n \xi_j R_j(t)$ where $\xi_0 = \eta_0$.

The key points are (i) the elements R_i that allow us to express a function evaluated at a point as an inner-product constraint, and (ii) the restriction to functions in F . F is a very special function space, a reproducing kernel Hilbert space, and R_i is determined by a

reproducing kernel R .

Ultimately, our goal is to address more complicated regression problems like the

Linear Smoothing Spline Problem. Consider simple regression data (x_i, y_i) with $0 \leq x_i \leq 1$, $i = 1, \dots, n$, and finding the function that minimizes

$$\frac{1}{n} \sum_{i=1}^n \{y_i - f(x_i)\}^2 + \lambda \int_0^1 f'(x)^2 dx. \quad (11)$$


If $f(x)$ is restricted to be in some class of functions F , minimizing only the first term gives least squares estimation within F . If F contains functions with $f(x_i) = y_i$ for all i , such functions minimize the first term but are typically very “unsmooth,” i.e., have large second term. The second “penalty” term is minimized by having a horizontal line, but that rarely has a small first term. As we will see in Section 4, for suitable F the minimizer takes the form

$$\hat{f}(x) = \xi_0 + \sum_{i=1}^n \xi_i R_i(x),$$

where the R_i ’s are known functions and the ξ_i ’s are coefficients found by solving a system of linear equations. This produces a linear smoothing spline.

If our goal is only to derive the solution to the linear smoothing spline problem with one predictor variable, RKHS theory is overkill. The value of RKHS theory lies in its generality. The linear spline penalty can be replaced by any other penalty with an associated inner product, and the x_i ’s can be vectors in \mathfrak{R}^p . Using RKHS results, we can solve the general problem of finding the minimizer of $\frac{1}{n} \sum_{i=1}^n (y_i - f(\mathbf{x}_i))^2 + \lambda J(f)$ for a general functional J that corresponds to a squared norm in a subspace. See Wahba (1990) or Gu (2002) for a full treatment of this approach. We now present an introduction to this theory.

Figure 1: Linear Interpolating Spline



2 Vector, Banach and Hilbert Spaces

This section presents background material required for the formal development of the RKHS framework.

2.1 Vector Spaces

A vector space is a set that contains elements called “vectors” and supports two kinds of operations: addition of vectors and multiplication by scalars. The scalars are drawn from some field (the real numbers in the rest of this article) and the vector space is said to be a vector space over that field. Formally, a set V is a vector space over a field F if there exists a structure $\{V, F, +, \times, 0_v\}$ consisting of V , F , a vector addition operation $+$, a scalar multiplication \times , and an identity element $0_v \in V$. This structure must obey the following axioms for any $\mathbf{u}, \mathbf{v}, \mathbf{w} \in V$ and $a, b \in F$:

- Associative Law: $(\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$.
- Commutative Law: $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$.
- Inverse Law: $\exists \mathbf{s} \in V$ s.t. $\mathbf{u} + \mathbf{s} = 0_v$. (Write $-\mathbf{u} \equiv \mathbf{s}$.)
- Identity Laws:
 - $0_v + \mathbf{u} = \mathbf{u}$.
 - $1 \times \mathbf{u} = \mathbf{u}$.
- Distributive Laws:
 - $a \times (b \times \mathbf{u}) = (a \times b) \times \mathbf{u}$.
 - $(a + b) \times \mathbf{u} = a \times \mathbf{u} + b \times \mathbf{u}$.
 - $a \times (\mathbf{u} + \mathbf{v}) = a \times \mathbf{u} + a \times \mathbf{v}$.

We will write $\mathbf{0}$ for $0_v \in V$ and $\mathbf{u} + (-\mathbf{v})$ as $\mathbf{u} - \mathbf{v}$. Any subset of a vector space that is closed under vector addition and scalar multiplication is called a subspace.

The simplest example of a vector space is just \mathfrak{R} itself, which is a vector space over \mathfrak{R} . Vector addition and scalar multiplication are just addition and multiplication on \mathfrak{R} . For more on vector spaces and the other topics to follow in this section, see, for example, Naylor & Sell (1982), Young (1988), Máté (1990) and Rustagi (1994).

2.2 Banach Spaces

A Banach space has a level of additional structure over that required to be a vector space. It is a vector space that also has a distance measure called a “norm” and is “complete” under that norm. Defining a Banach space sets the stage for defining a Hilbert space, which involves an additional bit of structure (an “inner product”) beyond that required to be a Banach space.

Definition. A norm of a vector space V , denoted by $\|\cdot\|$, is a nonnegative real valued function satisfying the following properties for all $\mathbf{u}, \mathbf{v} \in V$ and all $a \in \mathfrak{R}$,

1. Non-negative: $\|\mathbf{u}\| \geq 0$
2. Strictly positive: $\|\mathbf{u}\| = 0$ implies $\mathbf{u} = 0$
3. Homogeneous: $\|a\mathbf{u}\| = |a| \|\mathbf{u}\|$
4. Triangle inequality: $\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$

Definition. A vector space is called a normed vector space when a norm is defined on the space.

Definition. A sequence $\{\mathbf{v}_n\}$ in a normed vector space V is said to converge to $\mathbf{v}_0 \in V$ if

$$\lim_{n \rightarrow \infty} \|\mathbf{v}_n - \mathbf{v}_0\| = 0$$

Definition. A sequence $\{\mathbf{v}_n\} \subset V$ is called a Cauchy sequence if for any given $\epsilon > 0$, there exists an integer N such that

$$\|\mathbf{v}_m - \mathbf{v}_n\| < \epsilon, \quad \text{whenever } m, n \geq N.$$

Convergence of sequences in normed vector spaces follows the same general idea as sequences of real numbers except that the distance between two elements of the space is measured by the norm of the difference between the two elements.

Definition (Banach Space). A normed vector space V is called complete if every Cauchy sequence in V converges to an element of V . A complete normed vector space is called a Banach Space.

Example 2.1. \mathfrak{R} with the absolute value norm $\|x\| \equiv |x|$ is a complete, normed vector space over \mathfrak{R} , and is thus a Banach space.

Example 2.2. Let $\mathbf{x} = (x_1, \dots, x_n)^t$ be a point in \mathfrak{R}^n . The l_p norm on \mathfrak{R}^n is defined by

$$\|\mathbf{x}\|_p = \left[\sum_{i=1}^n |x_i|^p \right]^{1/p} \quad \text{for} \quad 1 \leq p < \infty.$$

One can verify properties 1-4 at the beginning of this section for each p , validating that $\|\mathbf{x}\|_p$ is a norm on \mathfrak{R}^n . Under the l_p norm, \mathfrak{R}^n is complete and thus a Banach space.

2.3 Hilbert Spaces

A Hilbert Space is a Banach space in which the norm is defined by an inner-product (also called dot-product) which we define below. We typically denote Hilbert spaces by H . For elements $\mathbf{u}, \mathbf{v} \in H$, write the inner product of \mathbf{u} and \mathbf{v} either as $\langle \mathbf{u}, \mathbf{v} \rangle_H$ or, when it is clear by context that the inner product is taking place in H , as $\langle \mathbf{u}, \mathbf{v} \rangle$. If H is a vector space over F , the result of the inner product is an element in F . We have $F = \mathfrak{R}$, so the result of an inner product will be a real number. The inner product operation must satisfy four properties for all $\mathbf{u}, \mathbf{v}, \mathbf{w} \in H$ and all $a \in F$.

1. Associative: $\langle a\mathbf{u}, \mathbf{v} \rangle = a\langle \mathbf{u}, \mathbf{v} \rangle$.
2. Commutative: $\langle \mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{v}, \mathbf{u} \rangle$.
3. Distributive: $\langle \mathbf{u}, \mathbf{v} + \mathbf{w} \rangle = \langle \mathbf{u}, \mathbf{v} \rangle + \langle \mathbf{u}, \mathbf{w} \rangle$.

4. Positive Definite: $\langle \mathbf{u}, \mathbf{u} \rangle \geq 0$ with equality holding only if $\mathbf{u} = \mathbf{0}$.

Definition. A vector space with an inner product defined on it is called an inner-product space. The norm of an element u in an inner-product space is taken as $\|\mathbf{u}\| = \langle \mathbf{u}, \mathbf{u} \rangle^{1/2}$. Two vectors are said to be orthogonal if their inner product is 0 and two sets of vectors are said to be orthogonal if every vector in one is orthogonal to every vector in the other. The set of all vectors orthogonal to a subspace is called the orthogonal complement of the subspace. A complete inner-product space is called a Hilbert space.

Example 2.3. \mathbb{R}^n with inner product defined by

$$\langle \mathbf{u}, \mathbf{v} \rangle \equiv \mathbf{u}^t \mathbf{v} = \sum_{i=1}^n u_i v_i$$

is a Hilbert space. For any positive definite matrix \mathbf{A} , $\langle \mathbf{u}, \mathbf{v} \rangle \equiv \mathbf{u}^t \mathbf{A} \mathbf{v}$ also defines a valid inner product.

Example 2.4. Let $L^2(a, b)$ be the vector space of all real valued functions defined on the interval (a, b) that are square integrable and define the inner product

$$\langle f, g \rangle \equiv \int_a^b f(x)g(x)dx.$$

The inner-product space $L^2(a, b)$ is well-known to be complete, see de Barra (1981), thus $L^2(a, b)$ is a Hilbert space.

3 Reproducing Kernel Hilbert Spaces (RKHSs)

Hilbert spaces that display certain properties on certain linear operators are called reproducing kernel Hilbert spaces.

Definition. A function T mapping a vector space X into another vector space Y is called a linear operator if $T(\lambda_1 \mathbf{x}_1 + \lambda_2 \mathbf{x}_2) = \lambda_1 T(\mathbf{x}_1) + \lambda_2 T(\mathbf{x}_2)$ for any $\mathbf{x}_1, \mathbf{x}_2 \in X$ and any $\lambda_1, \lambda_2 \in \mathbb{R}$.

Any $m \times n$ matrix \mathbf{A} maps vectors in \mathfrak{R}^n into vectors in \mathfrak{R}^m via $\mathbf{A}\mathbf{x} = \mathbf{y}$ and is linear.

Definition. The operator $T : X \rightarrow Y$ mapping a Banach space into a Banach space is continuous at $\mathbf{x}_0 \in X$ if and only if for every $\epsilon > 0$ there exists $\delta = \delta(\epsilon) > 0$ such that for every \mathbf{x} with $\|\mathbf{x} - \mathbf{x}_0\| < \delta$ we have $\|T\mathbf{x} - T\mathbf{x}_0\| < \epsilon$.

Linear operators are continuous everywhere if they are continuous at 0.

Definition. A real valued function defined on a vector space is called a functional.

A $1 \times n$ matrix defines a linear functional on \mathfrak{R}^n .

Example 3.1. Let S be the set of bounded real valued continuous functions $\{f(x)\}$ defined on the real line. Then S is a vector space with the usual $+$ and \times operations for functions. Some functionals on S are $\phi(f) = \int_a^b f(x)dx$ and $\phi_a(f) = f'(a)$ for some fixed a and b . A functional of particular importance is the evaluation functional.

Definition. Let V be a vector space of functions defined from E into \mathfrak{R} . For any $t \in E$, denote by e_t the evaluation functional at the point t , i.e., for $g \in V$, the mapping is $e_t(g) = g(t)$.

For $V = \mathfrak{R}^p$, vectors can be viewed as functions from the set $E = \{1, 2, \dots, p\}$ into \mathfrak{R} . An evaluation functional is $e_i(\mathbf{x}) = x_i$. Clearly, evaluation functionals are linear operators.

In a Hilbert space (or any normed vector space) of functions, the notion of pointwise convergence is related to the continuity of the evaluation functionals. The following are equivalent for a normed vector space H of real valued functions.

- (i) The evaluation functionals are continuous for all $t \in E$.
- (ii) If $f_n, f \in H$ and $\|f_n - f\| \rightarrow 0$ then $f_n(t) \rightarrow f(t)$ for every $t \in E$.
- (iii) For every $t \in E$ there exists $K_t > 0$ such that $|f(t)| \leq K_t \|f\|$ for all $f \in H$.

Here (ii) is the definition of (i). See Máté (1990), pg. 123 for a proof of (iii).

To define a reproducing kernel, we need the famous *Riesz Representation Theorem*.

Theorem. Let H be a Hilbert space and let ϕ be a continuous linear functional on H . Then there is one and only one vector $g \in H$ such that

$$\phi(f) = \langle f, g \rangle, \quad \text{for all } f \in H.$$

The vector g is sometimes called the *representation* of ϕ . However, ϕ and g are different objects: ϕ is a linear functional on H and g is a vector in H . For a proof of this theorem see Naylor & Sell (1982) or Máté (1990), pg. 84.

Recall, for $H = \Re^p$, an evaluation functional is $e_i(\mathbf{x}) = x_i$. The representation of this linear functional is the indicator vector \mathbf{e}_i that is 0 everywhere except has a 1 in the i th place. Then

$$x_i = e_i(\mathbf{x}) = \mathbf{x}^t \mathbf{e}_i.$$

In fact, the entire representation theorem is well known in \Re^p because for $\phi(\mathbf{x})$ to be a linear functional there must exist a vector $\boldsymbol{\phi}$ such that

$$\phi(\mathbf{x}) = \boldsymbol{\phi}^t \mathbf{x}.$$

An element of a set of functions, say f , is sometimes denoted $f(\cdot)$ to be explicit that the elements are functions, whereas $f(t)$ is the value of $f(\cdot)$ evaluated at $t \in E$. Applying the Riesz Representation Theorem to a Hilbert space H of real valued functions in which evaluation functionals are continuous, for every $t \in E$ there is a unique symmetric function $R : E \times E \rightarrow \Re$ with $R(\cdot, t) \in H$ the representation of e_t , so that

$$f(t) = e_t(f) = \langle f(\cdot), R(\cdot, t) \rangle_H, \quad f \in H.$$

The function R is called a *reproducing kernel* (r.k.) and $f(t) = \langle f(\cdot), R(\cdot, t) \rangle$ is called the *reproducing property* of R . In particular, by the reproducing property

$$R(s, t) = \langle R(\cdot, t), R(\cdot, s) \rangle.$$

In Subsection 1.2 we found the r.k. for the linear interpolating spline problem, $R(s, t) = \min(s, t)$. For any fixed t , $R(\cdot, t)$ is a part of the space F defined there, since $R(0, t) = 0$,

and $\int_0^1 [\partial R(s, t) / \partial s]^2 ds < \infty$. Also, R has the reproducing property since

$$\begin{aligned} \langle f, R(\cdot, t) \rangle &= \int_0^1 [f'(s) \partial R(s, t) / \partial s] ds \\ &= \int_0^t f'(s) 1 ds + \int_t^1 f'(s) 0 ds \\ &= f(t) - f(0) = f(t). \end{aligned}$$

Many other, more detailed, examples involving RKHSs and their reproducing kernels follow, but first the formal definition of an RKHS is presented below.

Definition A Hilbert space H of functions defined on E is called a reproducing kernel Hilbert space if all evaluation functionals are continuous.

We now present several RKHS examples that we will then use to solve some familiar penalized regression problems in Section 4.

3.1 Examples of Reproducing Kernel Hilbert Spaces

Example 3.2. Consider the space of all constant functionals over $\mathbf{x} = (x_1, x_2, \dots, x_p)^t \in \mathbb{R}^p$,

$$H = \{f_\theta : f_\theta(\mathbf{x}) = \theta, \theta \in \mathbb{R}\},$$

with $\langle f_\theta, f_\lambda \rangle = \theta\lambda$. (For simplicity, think of $p = 1$.) Since \mathbb{R}^p is a Hilbert Space, so is H . H has continuous evaluation functionals, so it is an RKHS and has a unique reproducing kernel. To find the r.k., observe that $R(\cdot, \mathbf{x}) \in H$, so it is a constant for any \mathbf{x} . Write $R(\mathbf{x}) \equiv R(\cdot, \mathbf{x})$. By the representation theorem and the defined inner product

$$\theta = f_\theta(\mathbf{x}) = \langle f_\theta(\cdot), R(\cdot, \mathbf{x}) \rangle = \theta R(\mathbf{x})$$

for any \mathbf{x} and θ . This implies that $R(\mathbf{x}) \equiv 1$ so that $R(\cdot, \mathbf{x}) = R(\mathbf{x}) \equiv 1$ and $R(\cdot, \cdot) \equiv 1$.

Example 3.3. Consider all linear functionals over $\mathbf{x} \in \mathbb{R}^p$ passing through the origin

$$H = \{f_\theta : f_\theta(\mathbf{x}) = \theta^t \mathbf{x}, \theta \in \mathbb{R}^p\}.$$

Define $\langle f_\theta, f_\lambda \rangle = \theta^t \lambda = \theta_1 \lambda_1 + \theta_2 \lambda_2 + \dots + \theta_p \lambda_p$. The kernel R must satisfy

$$f_\theta(\mathbf{x}) = \langle f_\theta(\cdot), R(\cdot, \mathbf{x}) \rangle$$

for all θ and any \mathbf{x} . Since $R(\cdot, \mathbf{x}) \in H$, $R(\mathbf{v}, \mathbf{x}) = \mathbf{u}^t \mathbf{v}$ for some \mathbf{u} that depends on \mathbf{x} , i.e., $R(\cdot, \mathbf{x}) = f_{\mathbf{u}(\mathbf{x})}(\cdot)$, so $R(\mathbf{v}, \mathbf{x}) = \mathbf{u}(\mathbf{x})^t \mathbf{v}$. By our definition of H we have

$$\theta^t \mathbf{x} = f_\theta(\mathbf{x}) = \langle f_\theta(\cdot), R(\cdot, \mathbf{x}) \rangle = \langle f_\theta(\cdot), f_{\mathbf{u}(\mathbf{x})}(\cdot) \rangle = \theta^t \mathbf{u}(\mathbf{x}),$$

so we need $\mathbf{u}(\mathbf{x})$ such that for any θ and \mathbf{x} we have

$$\theta^t \mathbf{x} = \theta^t \mathbf{u}(\mathbf{x}).$$

It follows that $\mathbf{u}(\mathbf{x}) = \mathbf{x}$. For example, taking θ to be the indicator vector \mathbf{e}_i implies that $u_i(\mathbf{x}) = x_i$ for every $i = 1, \dots, p$. We now have $R(\cdot, \mathbf{x}) = f_{\mathbf{x}}(\cdot)$ so that

$$R(\tilde{\mathbf{x}}, \mathbf{x}) = \mathbf{x}^t \tilde{\mathbf{x}} = x_1 \tilde{x}_1 + x_2 \tilde{x}_2 + \dots + x_p \tilde{x}_p.$$

Before moving on to the next example, we present one further concept that is useful in RKHS approaches to regression problems.

The projection principle for a RKHS.

Consider the connection between the reproducing kernel R of the RKHS H and the reproducing kernel R_0 for a subspace $H_0 \subset H$. Let H_0^\perp be the orthogonal complement of H_0 . Then any vector $f \in H$ can be written uniquely as $f = f_0 + f_1$ with $f_0 \in H_0$ and $f_1 \in H_0^\perp$. More particularly, $R(\cdot, t) = R_0(\cdot, t) + R_1(\cdot, t)$ with $R_0(\cdot, t) \in H_0$ and $R_1(\cdot, t) \in H_0^\perp$ if and only if R_0 is the r.k. of H_0 and R_1 is the r.k. of H_0^\perp . For a proof see Gu (2002).

Example 3.4. Now consider all affine (i.e., linear plus a constant) functionals in \Re^p ,

$$H = \{f_\theta : f_\theta(\mathbf{x}) = \theta_0 + \theta_1 x_1 + \dots + \theta_p x_p, \theta \in \Re^{p+1}\},$$

with $\langle f_\theta, f_\lambda \rangle = \theta_0 \lambda_0 + \theta_1 \lambda_1 + \dots + \theta_p \lambda_p$. The subspace $H_0 = \{f_\theta \in H : \theta_0 \in \Re, 0 = \theta_1 = \dots = \theta_p\}$ has the orthogonal complement $H_0^\perp = \{f_\theta \in H : 0 = \theta_0\}$. For practical purposes, H_0 is

the space of constant functionals from Example 3.2 and H_0^\perp is the space of linear functionals from Example 3.3. Note that the inner product on H when applied to vectors in H_0 and H_0^\perp , respectively, reduces to the inner products used in Examples 3.2 and 3.3.

Write H as $H = H_0 \oplus H_0^\perp$ where \oplus denotes the *direct sum* of two vector spaces. For two subspaces A and B contained in a vector space C , the direct sum is the space $D = \{a + b : a \in A, b \in B\}$. Any elements $d_1, d_2 \in D$ can be written as $a_1 + b_1$ and $a_2 + b_2$, respectively for some $a_1, a_2 \in A$ and $b_1, b_2 \in B$. When the two subspaces are orthogonal, as in our example, those decompositions are unique and the inner product between d_1 and d_2 is $\langle d_1, d_2 \rangle = \langle a_1, a_2 \rangle + \langle b_1, b_2 \rangle$. For more information about direct sum decomposition, see Berlinet & Thomas-Agnan (2004) or Gu (2002), for example.

We have already derived the r.k.'s for H_0 and H_0^\perp (call them R_0 and R_1 , respectively) in Examples 3.2 and 3.3. Applying the projection principle, the r.k. for H is the sum of R_0 and R_1 , i.e.,

$$R(\tilde{\mathbf{x}}, \mathbf{x}) = 1 + \mathbf{x}^t \tilde{\mathbf{x}}.$$

Example 3.5. Denote by V the collection of functions f with $f'' \in L^2[0, 1]$ and consider the subspace

$$W_2^0 = \{f(x) \in V : f, f' \text{ absolutely continuous and } f(0) = f'(0) = 0\}.$$

Define an inner product on W_2^0 as

$$\langle f, g \rangle = \int_0^1 f''(t)g''(t)dt. \quad (12)$$

Below we demonstrate that for $f \in W_2^0$ and any s , $f(s)$ can be written as

$$f(s) = \int_0^1 (s - u)_+ f''(u) du, \quad (13)$$

where $(a)_+$ is a for $a > 0$ and 0 for $a \leq 0$. Given any arbitrary and fixed $s \in [0, 1]$,

$$\int_0^1 (s - u)_+ f''(u) du = \int_0^s (s - u) f''(u) du.$$

Integrating by parts

$$\int_0^s (s-u)f''(u)du = (s-s)f'(s) - (s-0)f'(0) + \int_0^s f'(u)du = \int_0^s f'(u)du$$

and applying the Fundamental Theorem of Calculus to the last term,

$$\int_0^s (s-u)f''(u)du = f(s) - f(0) = f(s) .$$

Since the r.k. of the space W_2^0 must satisfy $f(s) = \langle f(\cdot), R(\cdot, s) \rangle$ from (12) and (13) we see that $R(\cdot, s)$ is a function such that

$$\frac{d^2 R(u, s)}{du^2} = (s-u)_+.$$

We also know that $R(\cdot, s) \in W_2^0$, so using $R(s, t) = \langle R(\cdot, t), R(\cdot, s) \rangle$

$$R(s, t) = \int_0^1 (t-u)_+(s-u)_+ du = \frac{\max(s, t) \min^2(s, t)}{2} - \frac{\min^3(s, t)}{6}.$$

For further examples of RKHSs with various inner products, see Berlinet & Thomas-Agnan (2004).

4 Penalized Regression with RKHSs

As mentioned in the introduction, nonparametric regression is a powerful approach for solving many current problems. The nonparametric regression model is given by

$$y_i = f(\mathbf{x}_i) + \epsilon_i, \quad i = 1, 2, \dots, n,$$

where f is an unknown regression function and the ϵ_i are independent error terms. We start this section with two common examples of penalized regression: ridge regression and smoothing splines.

Ridge Regression. In the classical linear regression setting $y_i = \mathbf{x}_i^t \boldsymbol{\beta} + \epsilon_i$ the ridge regression estimator $\hat{\boldsymbol{\beta}}_R$ proposed by Hoerl & Kennard (1970) minimizes

$$\frac{1}{n} \sum_{i=1}^n \left(y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij} \right)^2 + \lambda \sum_{j=1}^p \beta_j^2 \quad (14)$$

where x_{ij} is the i th observation of the j th component. The resulting estimate is biased but can reduce the variance relative to least squares estimates. The tuning parameter $\lambda \geq 0$ is a constant that controls the trade-off between bias and variance in $\hat{\beta}_R$, and is often selected by some form of cross validation; see Section 4.4.

Smoothing Splines. Smoothing splines are among the most popular methods for the estimation of f , due to their good empirical performance and sound theoretical support. It is often assumed, without loss of generality, that the domain of f is $[0, 1]$. With $f^{(m)}$ the m -th derivative of f , a smoothing spline estimate \hat{f} is the unique minimizer of

$$\frac{1}{n} \sum_{i=1}^n \{y_i - f(x_i)\}^2 + \lambda \int [f^{(m)}(x)]^2 dx \quad (15)$$

The minimization of (15) is implicitly over functions with square integrable m -th derivatives. The first term of (15) encourages the fitted f to be close to the data, while the second term penalizes the roughness of f . The smoothing parameter λ , usually pre-specified, controls the trade-off between the two conflicting goals. The special case of $m = 1$ reduces to the linear smoothing spline problem from (11). In practice it is common to choose $m = 2$ in which case the minimizer f_λ of (15) is called a cubic smoothing spline. As $\lambda \rightarrow \infty$, \hat{f}_λ approaches the least squares simple linear regression line, while as $\lambda \rightarrow 0$, \hat{f}_λ approaches the minimum curvature interpolant.

4.1 Solving the General Penalized Regression Problem

We now review a general framework to minimize (14), (15) and many other similar criteria, cf. (O’Sullivan, Yandell & Raynor 1986, Lin & Zhang 2006, Storlie, Bondell, Reich & Zhang 2010, Storlie, Bondell & Reich 2010, Gu & Qiu 1993). The data model is

$$y_i = f(\mathbf{x}_i) + \epsilon_i, \quad i = 1, 2, \dots, n, \quad (16)$$

where the ϵ_i are error terms and $f \in V$, a given vector space of functions on a set E .

An estimate of f is obtained by minimizing

$$\frac{1}{n} \sum_{i=1}^n \{y_i - f(\mathbf{x}_i)\}^2 + \lambda J(f), \quad (17)$$

over $f \in V$ where J is a penalty functional that must satisfy several restrictions and that helps to define the vector space V . We require (i) that $J(f) \geq 0$ for any f in some vector space \tilde{V} , (ii) that the null set $N = \{f \in \tilde{V} : J(f) = 0\}$ be a subspace, i.e., that it be closed under vector addition and scalar multiplication, (iii) that for $f_N \in N$ and $f \in \tilde{V}$, we have $J(f_N + f) = J(f)$, and (iv) that there exists an RKHS H contained in \tilde{V} for which the inner product satisfies $\langle f, f \rangle = J(f)$. This condition forces the intersection of N and H to contain only the zero vector.

Define a finite dimensional subspace of N , say N_0 , with a basis of known functions, say $\{\phi_1, \dots, \phi_M\}$, and $M \leq n$. In applications N is often finite dimensional, and we can simply take $N_0 = N$. In the minimization problem, we restrict attention to $f \in V$ where

$$V \equiv N_0 \oplus H.$$

For Example 3.5, \tilde{V} consists of functions in $L^2[0, 1]$ with finite values of

$$J(f) \equiv \int_0^1 [f''(t)]^2 dt.$$

$J(f)$ satisfies our four conditions with $H = W_2^0$. The linear functions $f(x) = a + bx$ are in N so we can take $\phi_1(x) \equiv 1$ and $\phi_2(x) = x$. Note that equation (12) defines an inner product on W_2^0 but does not define an inner product on all of \tilde{V} because nonzero functions could have a zero inner product with themselves, hence the nontrivial nature of N .

The key result (Wahba's Representation Theorem, also known as the "dual form" or "kernel trick" (Pearce & Wand 2006)) is that the minimizer of (17) is a linear combination of known functions involving the reproducing kernel on H . This fact will allow us to find the coefficients of the linear combination by solving a quadratic minimization problem similar

to those in standard linear models.

Representation Theorem. The minimizer \hat{f}_λ of equation (17) has the form

$$\hat{f}_\lambda(\mathbf{x}) = \sum_{j=1}^M d_j \phi_j(\mathbf{x}) + \sum_{i=1}^n c_i R(\mathbf{x}_i, \mathbf{x}), \quad (18)$$

where $R(\mathbf{s}, \mathbf{t})$ is the r.k. for H . An informal proof is given below, see Wahba (1990) or Gu (2002) for a formal proof.

Since we are working in V , clearly, a minimizer \hat{f} must have $\hat{f} = \hat{f}_0 + \hat{f}_1$ with $\hat{f}_0 \in N_0$ and $\hat{f}_1 \in H$. We want to show that $\hat{f}_1(\cdot) = \sum_{i=1}^n c_i R(\mathbf{x}_i, \cdot)$. To simplify notation write

$$\hat{f}_R(\cdot) \equiv \sum_{i=1}^n c_i R(\mathbf{x}_i, \cdot). \quad (19)$$

Decompose H as $H = H_0 \oplus H_0^\perp$ where $H_0 = \text{span}\{R(\mathbf{x}_i, \cdot), i = 1, \dots, n\}$ so that

$$\hat{f}_1(\cdot) = \hat{f}_R(\cdot) + \eta(\cdot),$$

with $\eta(\cdot) \in H_0^\perp$. By orthogonality and the reproducing property of the r.k.,

$$0 = \langle R(\mathbf{x}_i, \cdot), \eta(\cdot) \rangle = \eta(\mathbf{x}_i).$$

We now establish the representation theorem. Using our assumptions about J ,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \{y_i - \hat{f}(\mathbf{x}_i)\}^2 + \lambda J(\hat{f}) &= \frac{1}{n} \sum_{i=1}^n \{y_i - \hat{f}_0(\mathbf{x}_i) - \hat{f}_1(\mathbf{x}_i)\}^2 + \lambda J(\hat{f}_0 + \hat{f}_1) \\ &= \frac{1}{n} \sum_{i=1}^n \{y_i - \hat{f}_0(\mathbf{x}_i) - \hat{f}_1(\mathbf{x}_i)\}^2 + \lambda J(\hat{f}_1) \\ &= \frac{1}{n} \sum_{i=1}^n \{y_i - \hat{f}_0(\mathbf{x}_i) - \hat{f}_R(\mathbf{x}_i) - \eta(\mathbf{x}_i)\}^2 + \lambda J(\hat{f}_R + \eta). \end{aligned}$$

Because $\eta(\mathbf{x}_i) = 0$ and using orthogonality within H

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \{y_i - \hat{f}(\mathbf{x}_i)\}^2 + \lambda J(\hat{f}) &= \frac{1}{n} \sum_{i=1}^n \{y_i - \hat{f}_0(\mathbf{x}_i) - \hat{f}_R(\mathbf{x}_i)\}^2 + \lambda \left[J(\hat{f}_R) + J(\eta) \right] \\ &\geq \frac{1}{n} \sum_{i=1}^n \{y_i - \hat{f}_0(\mathbf{x}_i) - \hat{f}_R(\mathbf{x}_i)\}^2 + \lambda J(\hat{f}_R). \end{aligned}$$

Clearly, any $\eta \neq 0$ makes the inequality strict, so minimizers have $\eta = 0$, $\hat{f} = \hat{f}_0 + \hat{f}_R$, and the last inequality an equality.

A remarkable feature of the result in (18) is that the form of the minimizer is represented by a finite dimensional basis, regardless of the dimension of H . For example, H could be all functions with second derivative in L^2 , such as in the cubic smoothing spline problem. This H would require an infinite expansion of basis functions to represent all functions in the space, yet the *solution* of the minimization can be represented by a finite basis! So once we know that the minimizer takes the form (18), we can find the coefficients of the linear combination by solving a quadratic minimization problem similar to those in standard linear models. This occurs because we can write $J(\hat{f}) = J(\hat{f}_R)$ as a quadratic form in $\mathbf{c} = (c_1, \dots, c_n)^t$. Define $\mathbf{\Sigma}$ as the $n \times n$ matrix where the i, j entry is $\Sigma_{ij} = R(\mathbf{x}_i, \mathbf{x}_j)$. The matrix $\mathbf{\Sigma}$ is commonly referred to as the *Gram* matrix (Wahba 1990, Gu 2002). Now, using the reproducing property of R , write

$$J(\hat{f}_R) = \left\langle \sum_{i=1}^n c_i R(\mathbf{x}_i, \cdot), \sum_{j=1}^n c_j R(\mathbf{x}_j, \cdot) \right\rangle = \sum_{i=1}^n \sum_{j=1}^n c_i c_j R(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{c}^t \mathbf{\Sigma} \mathbf{c}.$$

Define the observation vector $\mathbf{y} = [y_1, \dots, y_n]^t$, and let \mathbf{T} be the $n \times M$ matrix with the ij -th entry defined by $\mathbf{T}_{ij} = \{\phi_j(\mathbf{x}_i)\}$. The minimization of (17) then takes the form

$$\min_{\mathbf{c}, \mathbf{d}} \frac{1}{n} \|\mathbf{y} - (\mathbf{T}\mathbf{d} + \mathbf{\Sigma}\mathbf{c})\|^2 + \lambda \mathbf{c}^t \mathbf{\Sigma} \mathbf{c}. \quad (20)$$

To solve (20) we define the following matrices,

$$\mathbf{Q}_{n \times (n+M)} = \begin{bmatrix} \mathbf{T}_{n \times M} & \mathbf{\Sigma}_{n \times n} \end{bmatrix},$$

$$\mathbf{\gamma}_{(n+M) \times 1} = \begin{bmatrix} \mathbf{d}_{M \times 1} \\ \mathbf{c}_{n \times 1} \end{bmatrix},$$

and

$$\mathbf{S}_{(n+M) \times (n+M)} = \begin{bmatrix} \mathbf{0}_{M \times M} & \mathbf{0}_{M \times n} \\ \mathbf{0}_{n \times 1} & \mathbf{\Sigma}_{n \times n} \end{bmatrix}.$$

Now (20) becomes

$$\min_{\boldsymbol{\gamma}} \frac{1}{n} \|\mathbf{y} - \mathbf{Q}\boldsymbol{\gamma}\|^2 + \lambda \boldsymbol{\gamma}^t \mathbf{S} \boldsymbol{\gamma}. \quad (21)$$

The minimization in (21) is the same as a generalized ridge regression. Taking derivatives with respect to $\boldsymbol{\gamma}$ we have

$$(\mathbf{Q}^t \mathbf{Q} + \lambda \mathbf{S}) \hat{\boldsymbol{\gamma}} = \mathbf{Q}^t \mathbf{y},$$

which requires solving a system of $n + M$ equations to find $\hat{\boldsymbol{\gamma}}$. For analytical purposes, we can write $\hat{\boldsymbol{\gamma}}$ as

$$\hat{\boldsymbol{\gamma}} = (\mathbf{Q}^t \mathbf{Q} + \lambda \mathbf{S})^{-1} \mathbf{Q}^t \mathbf{y}. \quad (22)$$

As long as \mathbf{T} is of full column rank, then \hat{f}_λ is unique. If the \mathbf{x}_i are not unique, then $\hat{\boldsymbol{\gamma}}$ is not unique as defined, but \hat{f}_λ is still unique. One could simply use only the unique \mathbf{x}_i in the definition of f_R in (19) to ensure that $\hat{\boldsymbol{\gamma}}$ is unique in that case as well.

Alternatively, $\hat{\boldsymbol{\gamma}}$ can be obtained as the generalized least squares estimate from fitting the linear model

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{T} & \boldsymbol{\Sigma} \\ \mathbf{0}_{n \times M} & \mathbf{I}_{n \times n} \end{bmatrix} \begin{bmatrix} \mathbf{d} \\ \mathbf{c} \end{bmatrix} + e, \quad \text{Cov}(e) \propto \begin{bmatrix} \mathbf{I}_{n \times n} & \mathbf{0}_{n \times n} \\ \mathbf{0}_{n \times n} & (1/\lambda) \boldsymbol{\Sigma}_{n \times n}^{-1} \end{bmatrix}.$$

For clarity, we have restricted our attention to minimizing (17), which incorporates squared error loss between the observations and the unknown function evaluations. The representation theorem holds for more general loss functions (e.g., those from logistic or Poisson regression); see Gu (2002).

4.2 General Solution Applied to Cubic Smoothing Spline

Consider again the regression problem $y_i = f(x_i) + \epsilon_i, i = 1, 2, \dots, n$ where $x_i \in [0, 1]$ and $\epsilon_i \sim N(0, \sigma^2)$. We focus on the cubic smoothing spline solution to this problem. That is, we find a function that minimizes

$$\sum_{i=1}^n \{y_i - f(x_i)\}^2 + \lambda \int f''(x)^2 dx.$$

As discussed in Subsection 4.1, N is the space of linear functions from Example 3.4 with $p = 1$ (see also equation (24)) and a basis of $\phi_1(x) = 1$, $\phi_2(x) = x$. In this case, N happens to be an RKHS, but in general it is not even necessary to define an inner product on N . $H = W_2^0$ comes from Example 3.5. We know that the reproducing kernel for H is

$$R(s, t) = \int_0^1 (t - u)_+(s - u)_+ du = \frac{\max(s, t) \min^2(s, t)}{2} - \frac{\min^3(s, t)}{6},$$

so the solution has the form

$$\hat{f}(x) = \hat{d}_0(1) + \hat{d}_1 x + \sum_{i=1}^n \hat{c}_i R(x_i, x).$$

From (22) we have

$$(\mathbf{Q}^t \mathbf{Q} + \lambda \mathbf{S})^{-1} \mathbf{Q}^t \mathbf{y} = \begin{bmatrix} \hat{\mathbf{d}} \\ \hat{\mathbf{c}} \end{bmatrix} \quad (23)$$

The Supplementary Material A.1 provides code written in the R language (R Development Core Team 2005) and plots for fitting the cubic smoothing spline solution in (23) to some motorcycle accident data. The demonstration also includes searching for the best value of the tuning parameter λ which is briefly discussed in Section 4.4.

As an aside, it is well known that the basis functions $R(x_i, x)$ (that form the columns of \mathbf{Q} when evaluated at the data points) form a natural cubic spline with knots at the distinct values of x_i . See Wahba (1990) for a justification, which just involves some algebra. The $\max(x_i, x)$ and $\min(x_i, x)$ in $R(x_i, x)$ combine in a way to produce knots at the x_i , while the degree of the polynomial spline would clearly be three, since it is the highest power present in $R(s, t)$. This is the reason that the minimization problem in this section has been given the name “cubic” smoothing “spline”.

4.3 General Solution Applied to Ridge Regression

We now solve the linear ridge regression problem of minimizing (14) with the RKHS approach detailed above. Although the RKHS framework is not necessary to solve the ridge regression

problem, it serves as a good illustration of the RKHS machinery.

To put the ridge regression problem in the framework of (17) consider

$$\tilde{V} = \{f(\mathbf{x}) = \beta_0 + \sum_{j=1}^p \beta_j x_j\}. \quad (24)$$

with the penalty function

$$J(f) = \sum_{j=1}^p \beta_j^2.$$

Notice that by letting $\tilde{V} = V_0 \oplus H$ where V_0 is the RKHS from Example 3.2 and H is from Example 3.3, we have $J(f) = \langle f, f \rangle_H$. The dimension of $V_0 = N = N_0$ is $M = 1$ and $\phi_1(x) = 1$.

From (18) the solution takes the form

$$\hat{f}_\lambda(\mathbf{x}) = \hat{d}_1 + \sum_{i=1}^n \hat{c}_i R(\mathbf{x}_i, \mathbf{x}) \quad (25)$$

with

$$(\mathbf{Q}^t \mathbf{Q} + \lambda \mathbf{S})^{-1} \mathbf{Q}^t \mathbf{y} = \begin{bmatrix} \hat{d}_1 \\ \hat{\mathbf{c}} \end{bmatrix} \quad (26)$$

as given by (22). In this case, it is more familiar to write the solution in what Pearce & Wand (2006) refer to as the “primal” form,

$$\hat{f}_\lambda(\mathbf{x}) = \hat{\beta}_0 + \hat{\beta}_1 x_1 + \hat{\beta}_2 x_2 + \dots + \hat{\beta}_p x_p.$$

This can be done by recalling from Example 3.3 that

$$R(\mathbf{x}_i, \mathbf{x}) = x_{i1}x_1 + x_{i2}x_2 + \dots + x_{ip}x_p.$$

Substituting into (25) gives

$$\hat{f}(\mathbf{x}) = \hat{d}_1 + \sum_{i=1}^n \hat{c}_i x_{i1}x_1 + \sum_{i=1}^n \hat{c}_i x_{i2}x_2 + \dots + \sum_{i=1}^n \hat{c}_i x_{ip}x_p,$$

which implies that

$$\hat{\beta}_0 = \hat{d}_1 \quad \text{and} \quad \hat{\beta}_j = \sum_{i=1}^n \hat{c}_i x_{ij} \quad (27)$$

for $j = 1, 2, \dots, p$. In the Supplementary Material A.2, we demonstrate this solution on Longley's (1967) employment data using the statistical software R. The RKHS solution applied to ridge regression is mostly for illustrative (not practical) purposes, as it requires a $(n+1) \times (n+1)$ matrix solve. This is of course less efficient than the standard approach unless $p > n$.

Both of the previous examples illustrate a typical case in which \tilde{V} is constructed as the direct sum of a vector space V_0 and a RKHS H whose intersection is the zero vector,

$$\tilde{V} = V_0 \oplus H.$$

Because the intersection is zero, any $f \in \tilde{V}$ can be written uniquely as $f = f_0 + f_1$ with $f_0 \in V_0$ and $f_1 \in H$. The importance of this is that if the penalty functional J happens to satisfy

$$J(f) = \langle f_1, f_1 \rangle_H,$$

then all of our assumptions about J hold immediately with $N = V_0$, and the solution in (22) can be applied.

As a further aside, when V_0 is itself a RKHS, then \tilde{V} is an RKHS under the inner product

$$\langle f, g \rangle_{\tilde{V}} \equiv \langle f_0, g_0 \rangle_{V_0} + \langle f_1, g_1 \rangle_H,$$

with reproducing kernel $R_{\tilde{V}} = R_{V_0} + R_H$. Note that V_0 and H are orthogonal under this inner product. This orthogonal decomposition of \tilde{V} is also closely related to additive models (Wood 2006) and more generally Smoothing Spline ANOVA models (Gu 2002), which also include tensor product splines as a special case. Thin plate splines (Wahba 1990) also fall nicely into the general RKHS framework.

4.4 Choosing the Degree of Smoothness

With the penalized regression procedures described above, the choice of the smoothing parameter λ is an important issue. There are many methods available for this task; e.g., visual inspection of the fit, m -fold cross-validation (Kohavi 1995), AIC/unbiased risk estimation, generalized maximum likelihood (Wahba 1990), generalized cross-validation (GCV) (Craven & Wahba 1979), etc. For those examples given in the Supplementary Material, we use the GCV approach, which works as follows for any generalized ridge regression solution, such as in (26). Suppose that an estimate admits the following closed-form expression:

$$\hat{\beta} = (\mathbf{Q}^t \mathbf{Q} + \lambda \mathbf{S})^{-1} \mathbf{Q}^t \mathbf{y}.$$

The GCV choice of λ for this generalized ridge estimate is the minimizer of

$$V(\lambda) = \frac{1}{n} \|(\mathbf{I} - \mathbf{A}(\lambda))\mathbf{y}\|^2 / \left[\frac{1}{n} \text{Trace}\{\mathbf{I} - \mathbf{A}(\lambda)\} \right]^2,$$

where

$$\mathbf{A}(\lambda) = \mathbf{Q}(\mathbf{Q}^t \mathbf{Q} + \lambda \mathbf{S})^{-1} \mathbf{Q}^t.$$

The goal of GCV is to find the optimal λ so that the resulting $\hat{\beta}$ has the smallest mean squared error. For more details about GCV and other methods of finding λ see Golub, Heath & Wahba (1979), Allen (1974), Wecker & Ansley (1983) and Wahba (1990).

5 Concluding Remarks

We have given several examples illustrating the utility of the RKHS approach to penalized regression problems. We reviewed the building blocks necessary to define an RKHS and presented several key results about these spaces. Finally, we used the results to illustrate estimation for cubic smoothing spline problems and ridge regression, providing transparent R code to enhance understanding.

The reader is now encouraged to explore some more advanced papers and texts that they now have the tools to access, and unlock the full potential of RKHS methods. Wahba (1990)

and Gu (2002) discuss smoothing spline ANOVA, which is a flexible modeling framework ranging from additive modeling on one extreme to full tensor product splines on the other. Gu (2002) also covers generalized models in the smoothing spline ANOVA framework. Pearce & Wand (2006) explore the intimate connection between RKHSs in penalized regression and the much used support vector machine. Lin & Zhang (2006) and Storlie, Bondell, Reich & Zhang (2010) use the RKHS framework to develop a smoothing spline version of the popular LASSO (Tibshirani 1996) and adaptive LASSO (Zou 2006), respectively. The RKHS framework can also be used to do spatially adaptive smoothing (Storlie, Bondell & Reich 2010). The flexibility and elegance of RKHS methods are remarkable.

6 Acknowledgments

The authors are grateful to the referees and the associate editor for their comments which led to a much improved version of the paper. Raymond Wong also provided helpful comments on an earlier version of the paper. Nosedal was partially supported by the National Council of Science and Technology of Mexico (CONACYT). Lee was partially supported by the National Science Foundation under grant DMS 1007520. Storlie’s work was funded by Los Alamos National Security, LLC (LANS), operator of the Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396 with the U.S. Department of Energy. This paper is published under LA-UR-12-10169.

References

- Allen, D. (1974), ‘The relationship between variable selection and data augmentation and a method for prediction’, *Technometrics* **16**, 125–127.
- Berlinet, A. & Thomas-Agnan, C. (2004), *Reproducing Kernel Hilbert Spaces in Probability and Statistics*, Kluwer Academic Publishers, Norwell, MA.
- Berman, M. (1994), ‘Automated smoothing of image and other regularly spaced data’, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **16**, 460–468.

- Bivand, R., Pebesma, E. & Gómez-Rubio, V. (2008), *Applied spatial data analysis with R: Use R*, Springer.
- Craven, P. & Wahba, G. (1979), ‘Smoothing noisy data with spline functions: estimating the correct degree of smoothing by the method of generalized cross-validation’, *Numerical Mathematics* **31**, 377–403.
- de Barra, G. (1981), *Measure Theory and Integration*, Horwood Publishing, West Sussex, UK.
- Eubank, R. (1999), *Nonparametric Regression and Spline Smoothing*, Marcel Dekker, New York, NY.
- Golub, G., Heath, M. & Wahba, G. (1979), ‘Generalized cross-validation as a method for choosing a good ridge parameter’, *Technometrics* **21**, 215–223.
- Gu, C. (2002), *Smoothing Spline ANOVA Models*, Springer-Verlag, New York, NY.
- Gu, C. & Qiu, C. (1993), ‘Smoothing spline density estimation: Theory’, *Annals of Statistics* **21**, 217–234.
- Hastie, T., Tibshirani, R. & Friedman, J. (2001), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer-Verlag, New York, NY.
- Hoerl, A. & Kennard, R. (1970), ‘Ridge regression: Biased estimation for nonorthogonal problems’, *Technometrics* **12**, 55–67.
- Kohavi, R. (1995), ‘A study of cross-validation and bootstrap for accuracy estimation and model selection’, *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence* **2**(12), 1137–1143.
- Lin, Y. & Zhang, H. (2006), ‘Component selection and smoothing in smoothing spline analysis of variance models’, *Annals of Statistics* **34**, 2272–2297.
- Longley, J. (1967), ‘An appraisal of least-squares programs from the point of view of the user’, *Journal of the American Statistical Association* **62**, 819–841.
- Máté, L. (1990), *Hilbert Space Methods in Science and Engineering*, Taylor & Francis, Oxford, UK.
- Naylor, A. & Sell, G. (1982), *Linear Operator Theory in Engineering and Science*, Springer, New York, NY.
- O’Sullivan, F., Yandell, B. & Raynor, W. (1986), ‘Automatic smoothing of regression functions in generalized linear models’, *Journal of the American Statistical Association* **81**, 96–103.
- Pearce, N. & Wand, M. (2006), ‘Penalised splines and reproducing kernel methods’, *American Statistician* **60**, 233–240.

- R Development Core Team (2005), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
URL: <http://www.R-project.org>
- Ramsay, J. & Silverman, B. (2005), *Functional Data Analysis*, New York, NY: Springer-Verlag.
- Rustagi, J. (1994), *Optimization Techniques in Statistics*, Academic Press, London, UK.
- Silverman, B. (1985), ‘Some aspects of the spline smoothing approach to non-parametric curve fitting’, *Journal of the Royal Statistical Society: Series B* **47**, 1–52.
- Storlie, C., Bondell, H. & Reich, B. (2010), ‘A locally adaptive penalty for estimation of functions with varying roughness’, *Journal of Computational and Graphical Statistics* **19**, 569–589.
- Storlie, C., Bondell, H., Reich, B. & Zhang, H. (2010), ‘Surface estimation, variable selection, and the nonparametric oracle property’, *Statistica Sinica* **21**, 679–705.
- Storlie, C., Swiler, L., Helton, J. & Sallaberry, C. (2009), ‘Implementation and evaluation of nonparametric regression procedures for sensitivity analysis of computationally demanding models’, *Reliability Engineering and System Safety* **94**, 1735–1763.
- Tibshirani, R. (1996), ‘Regression shrinkage and selection via the lasso’, *Journal of the Royal Statistical Society: Series B* **58**, 267–288.
- Wahba, G. (1990), *Spline Models for Observational Data*, vol. 59. SIAM. CBMS-NSF Regional Conference Series in Applied Mathematics.
- Wecker, W. & Ansley, C. (1983), ‘The signal extraction approach to nonlinear regression and spline smoothing’, *Journal of the American Statistical Association* **78**, 81–89.
- Wood, S. (2006), *Generalized additive models: an introduction with R*, CRC Press, Boca Raton, FL.
- Young, N. (1988), *An Introduction to Hilbert Space*, Cambridge University Press, Cambridge, UK.
- Zou, H. (2006), ‘The adaptive lasso and its oracle properties’, *Journal of the American Statistical Association* **101**(476), 1418–1429.

A Supplementary Material: Examples using R

A.1 RKHS solution applied to Cubic Smoothing Spline

We apply the cubic smoothing spline within the RKHS framework on the benchmark motorcycle accident data used in Silverman (1985). These data come from a computer simulation of motorcycle accidents. The response is a series of measurements of head acceleration over time in a simulated motorcycle accident used to test crash helmets.

The following code imports the data, creates a function to calculate the cubic smoothing spline estimate, finds the optimal GCV λ value and corresponding cubic smoothing spline estimator, and creates the GCV plot and fitted line plot in Figure 2.

```
rm(list=ls(all=TRUE))

#### Data ####
library(MASS)      # Package that contains the mcycle dataset
x<-mcycle$times
x<-(x-min(x))/(max(x)-min(x))  ## need to scale predictor to [0,1]
y<-mcycle$accel

#### Reproducing Kernel for <f,g>=int_0^1 f''(x)g''(x)dx #####
rk.1<-function(s,t){
return( (1/2)*min(s,t)^2*max(s,t) - (1/6)*(min(s,t))^3 )
}

#### Function to obtain the Gram Matrix ####
get.gram.1<-function(X){
n<-dim(X)[1] ;
Gram<-matrix(0,n,n) #initializes Gram array ;
#i=index for rows
#j=index for columns
Gram<-as.matrix(Gram) # Gram matrix ;
for (i in 1:n){
  for (j in 1:n){
    Gram[i,j]<-rk.1(X[i,],X[j,])
  }
}
return(Gram) }
```

```
#### Function to find a stable gen-inverse of a symmetric nonneg def ####
#### matrix to avoid numerical problems for small lambdas ####
gen.inv<-function(X,eps=1e-12){
eig.X<-eigen(X,symmetric=T) ;
P<-eig.X[[2]] ;
lambda<-eig.X[[1]] ;
ind<-lambda>eps ;
lambda[ind]<-1/lambda[ind] ;
lambda[!ind]<-0 ;
ans<-P%*%diag(lambda,nrow=length(lambda))%*%t(P) ;
return(ans) }
```

```
#### Cubic Smoothing Spline Function ####
smoothing.spline<-function(X,y,lambda){
X<-as.matrix(X)
Gram<-get.gram.1(X) #Gram matrix (nxn) ;
n<-dim(X)[1] # n=length of y ;
J<-matrix(1,n,1) # vector of ones dim ;
T<-cbind(J,X) # matrix with a basis for the null space of the penalty ;
Q<-cbind(T,Gram) # design matrix ;
m<-dim(T)[2] # dimension of the null space of the penalty ;
S<-matrix(0,n+m,n+m) #initialize S ;
S[(m+1):(n+m),(m+1):(n+m)]<-Gram #non-zero part of S ;
M<-(t(Q)%*%Q+lambda*S) ;
M.inv<-gen.inv(M) # gen-inverse of M ;
gamma.hat<-crossprod(M.inv,crossprod(Q,y)) ;
f.hat<-Q%*%gamma.hat ;
A<-Q%*%M.inv%*%t(Q) ;
tr.A<-sum(diag(A)) #trace of hat matrix ;
rss<-t(y-f.hat)%*%(y-f.hat) #residual sum of squares ;
gcv<-n*rss/(n-tr.A)^2 #obtain GCV score ;
return(list(f.hat=f.hat,gamma.hat=gamma.hat,gcv=gcv))
}
```

```
#### Find an optimal lambda using GCV... ####
log.lambda<-seq(-6,0,by=.1) ;
V<-rep(0,length(log.lambda)) ;
for (i in 1:length(log.lambda)){
V[i]<-smoothing.spline(x,y,10^log.lambda[i])$gcv #obtain GCV score ;
}
min.ind<-order(V)[1] # extract index of min(V) ;
opt.mod.1<-smoothing.spline(x,y,10^log.lambda[min.ind]) #fit optimal model ;
```



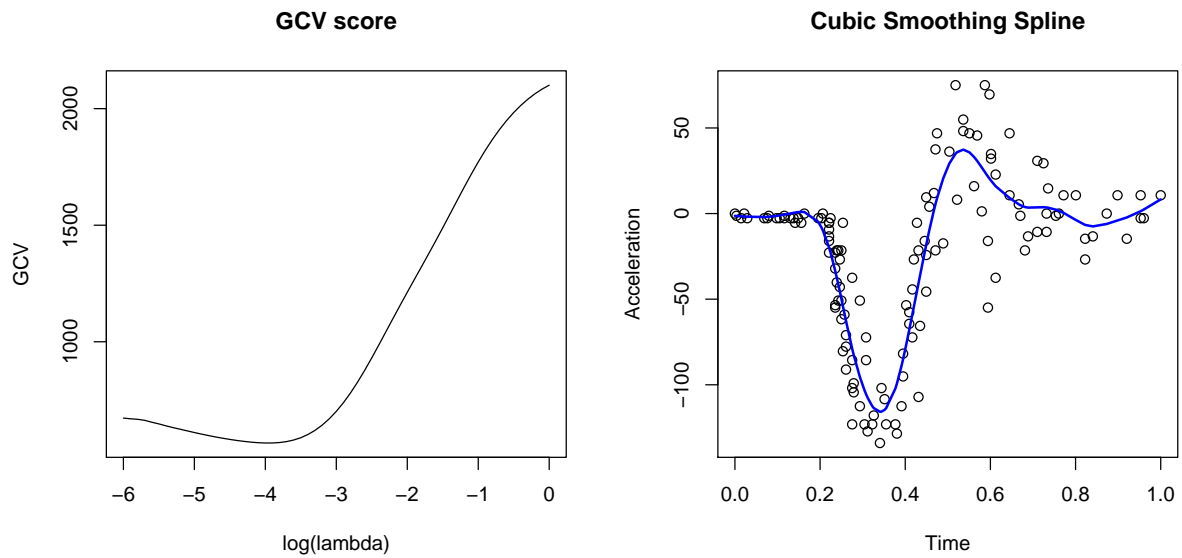
```

#### Plot of GCV ####
plot(log.lambda,V,type="l",main="GCV score",xlab="log(lambda)",ylab="GCV") ;

#### Fitted Line Plot for Cubic Smoothing Spline ####
plot(x,y,xlab="x",ylab="response",main="Cubic Smoothing Spline") ;
lines(x,opt.mod.1$f.hat,type="l",lty=1,lwd=2,col="blue") ;

```

Figure 2: Cubic Smoothing Spline GCV plot (left) and optimal fitted line plot (right)



A.2 RKHS solution applied to Ridge Regression

Here an example of ridge regression in the RKHS framework is provided using the labor statistics data of Longley (1967) which is known for having high collinearity. We seek to model Employment based on Prices, GNP, Unemployment, Military Size, Population Size, and Year. The following code, imports the data, creates a function to calculate the ridge regression estimate, finds the optimal GCV λ value and corresponding ridge regression estimator, creates the GCV plot in Figure 3, and finally converts the RKHS estimator back to slope coefficients as in (27) and compares to the ordinary least squares solution.

```

#### Data ####
data(longley)
y <- longley[,1] ;
X <- as.matrix(longley[,2:7]) ;
for(j in 1:ncol(X)) # scale each predictor to [0,1]
X[,j]<-(X[,j]-min(X[,j]))/(max(X[,j])-min(X[,j])) ;

#### Reproducing Kernel #####
rk.2<-function(s,t){
p<-length(s) ;
rk<-0 ;
for (i in 1:p){
rk<-s[i]*t[i]+rk
}
return( rk) )
}

#### Gram matrix #####
get.gram.2<-function(X){
n<-dim(X)[1];
Gram<-matrix(0,n,n);
#initializes Gram array;
#i=index for rows;
#j=index for columns;
for(i in 1:n){
  for (j in 1:n){
    Gram[i,j]<-rk.2(X[i,],X[j,])
  }
}
return(Gram) }

#### Ridge Regression Function ####
ridge.regression<-function(X,y,lambda){
Gram<-get.gram.2(X) #Gram matrix (nxn) ;
n<-dim(X)[1] # n=length of y ;
J<-matrix(1,n,1) # vector of ones dim ;
Q<-cbind(J,Gram) # design matrix ;
m<-1 # dimension of the null space of the penalty ;
S<-matrix(0,n+m,n+m) #initialize S ;
S[(m+1):(n+m),(m+1):(n+m)]<-Gram #non-zero part of S ;
M<-(t(Q)%*%Q+lambda*S) ;
M.inv<-gen.inv(M) # gen-inverse of M ;

```

```

gamma.hat<-as.numeric(crossprod(M.inv,crossprod(Q,y))) ;
f.hat<-Q%*%gamma.hat ;
A<-Q%*%M.inv%*%t(Q) ;
tr.A<-sum(diag(A)) #trace of hat matrix ;
rss<-t(y-f.hat)%*(y-f.hat) #residual sum of squares ;
gcv<-n*rss/(n-tr.A)^2 #obtain GCV score ;
beta.hat.0<-gamma.hat[1] #intercept ;
beta.hat<-c(beta.hat.0,gamma.hat[-1]*X) # slopes ;
return(list(f.hat=f.hat,gamma.hat=gamma.hat,gcv=gcv,beta.hat=beta.hat))
}

#### Find an optimal lambda using GCV... ####
log.lambda<-seq(-6,0,by=.1) ;
V<-rep(0,length(log.lambda)) ;
beta.mat<-matrix(0,length(log.lambda),ncol(X)+1) ;
for (i in 1:length(log.lambda)){
ans.i<-ridge.regression(X,y,10^log.lambda[i]) ;
V[i]<-ans.i$gcv #obtain GCV score ;
beta.mat[i,]<-ans.i$beta.hat #obtain beta estimates ;
}

#### Plot of GCV scores ####
plot(log.lambda,V,type="l",main="GCV score",xlab="log(lambda)",ylab="GCV") ;

#### Plot betas across lambda ####
plot(log.lambda,beta.mat[,2],main="Betas Across Lambda",xlab="log(lambda)",
ylab="Beta",col=0, ylim=range(beta.mat[,1])) ;
for(j in 1:ncol(X))
lines(log.lambda,beta.mat[,j+1],col=j) ;
legend(-2,80,legend=paste("Beta",1:ncol(X)),lty=rep(1,ncol(X)),col=1:ncol(X)) ;

### Fit GCV optimal ridge regression ####
min.ind<-order(V)[1] # extract index of min(V) ;
opt.mod.2<-ridge.regression(X,y,10^log.lambda[min.ind]) ;

#### Fit ordinary least squares regression (i.e., lambda=0, no penalty) ####
ols.mod.2<-ridge.regression(X,y,0) ;

#### Print Ridge Regression Estimates ####
round(opt.mod.2$beta.hat,2) ;
#### Print OLS Estimates ####
round(ols.mod.2$beta.hat,2) ;

```

The last two print statements (which include the intercept $\hat{\beta}_0$) give

$$\hat{\beta}_{\text{opt}} = [82.78, 54.17, 5.36, 1.38, -28.79, 5.40, -0.61]'$$

$$\hat{\beta}_{\text{ols}} = [81.64, 84.49, 10.71, 2.39, -39.04, -21.28, 2.40]'$$

Figure 3: Ridge Regression GCV plot (left) and change in $\hat{\beta}_j$ across λ (right)

